

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR UNITED STATES PATENT

FOR

DEVICE AND METHOD FOR DATA STORAGE

Inventors:

Stephen S. Jackson
104 Antler Point Rd
Chapel Hill, NC 27516

Ottis K. Black III
1201 Hillingdon Way
Raleigh, NC 27614

Attorney Docket: 2204/A81
(12974RNUS01U)

Attorneys:

BROMBERG & SUNSTEIN LLP
125 Summer Street
Boston, MA 02110
(617) 443-9292

002231" 45/84250

DEVICE AND METHOD FOR DATA STORAGE

FIELD OF THE INVENTION

5 The invention generally relates to communication networks and, more particularly, the invention relates to data storage devices that may receive power and data from a communication network.

BACKGROUND OF THE INVENTION

10 There are a number of approaches to providing data storage in networks (e.g., intranets and other local networks). A basic approach is to connect a storage device (e.g., a disk drive, RAID array, tape drive or optical drive) to another network device in the local area network, such as a server or personal computer, using an interface such as the Small Computer System Interface ("SCSI"). The SCSI interface allows network devices, such as a server, to
15 communicate with peripheral devices such as disk drives, tape drives, CD-ROM drives and printers. Figure 1A is a block diagram of a local area network 106 with a disk array 102 coupled to the local area network 106 through server 104. As shown in Figure 1A, the server 104 acts as a pathway between the disk array
20 102 and the local area network 106. The local area network 106 may be any type of local area network, such as a local area network that operates according to the Ethernet protocol.

 An alternative approach is a storage area network ("SAN") as shown in Figure 1B. A SAN architecture allows storage resources (*i.e.*, disk arrays 102) to
25 be shared among multiple network devices (*e.g.*, servers 104) in the local area network 106. The disk arrays 102 are attached to the servers 104 using, for example, switches 112. Typically, the disk arrays 102 are connected using a high-speed interconnection, such as Fibre Channel technology. The SAN architecture

provides any-to-any connectivity between the servers 104 and the disk arrays 102. As shown in Figure 1B, the disk arrays 102 are coupled to the local area network 106 via the switches 112 and the servers 104.

5

SUMMARY OF THE INVENTION

In accordance with one aspect of the invention, a data storage device for use in a power integrated network includes a memory configured for storing data and a control module coupled to the memory that controls the transmission of data from the memory to the power integrated network and the storage of data received from the power integrated network in the memory. In addition, a power module is coupled to the memory and the control module. The power module receives power from the power integrated network to energize the data storage device. The data storage device may further include a network interface coupled to the control module where the network interface is configured to communicate with the power integrated network. The control module may package data for transmission over the power integrated network.

In one embodiment, the power integrated network is a Power Ethernet network. The memory area may be a device including a SCSI interface. In another embodiment, the power module includes a power converter for converting the power received from a first voltage level to a second voltage level. The second voltage level may be lower than the first voltage level.

In accordance with another aspect of the invention, a method for providing data storage in a power integrated network includes coupling a data storage device to the power integrated network where the data storage device is configured to communicate with the power integrated network and receiving data and power at the data storage device from the power integrated network. The power received from the power integrated network is used to energize the data storage device. In one embodiment, the power integrated network is a

Power Ethernet. The method may also include converting the power received from the power integrated network from a first voltage level to a second voltage level. The second voltage level may be lower than the first voltage level.

In accordance with yet another aspect of the invention, a system for providing data storage in a communication network includes a first power integrated network and a second power integrated network coupled to a specified network so that data may be transmitted from and received by the first power integrated network and the second power integrated network across the specified network. The first power integrated network and the second power integrated network provide data and power. A first data storage device is coupled to the first power integrated network where the first data storage device is configured to communicate with the first power integrated network and to receive power from the first power integrated network. A second data storage device is coupled to the second power integrated network where the second data storage device is configured to communicate with the second power integrated network and to receive power from the second power integrated network.

In one embodiment, the first power integrated network and the second power integrated network are Power Ethernet networks. The specified network may be the Internet. In another embodiment, the first storage device and the second storage device are in a RAID configuration. In yet another embodiment, the first data storage device includes a memory configured to store data and a control module coupled to the memory area that controls the transmission of data from the memory to the first power integrated network and the storage of data received from the first power integrated network in the memory. The first data storage device further includes a power module coupled to the memory and the control module. The power module receives power from the first power integrated network to energize the first data storage device. In one embodiment, the memory may be a device including a SCSI interface.

In another embodiment, the first data storage device includes a network interface coupled to the control module where the network interface is configured to communicate with the first power integrated network. The control module may package data for transmission over the first power integrated network. In yet another embodiment, the power module includes a power converter for converting the power received from a first voltage level to a second voltage level. The second voltage level may be lower than the first voltage level.

In a further embodiment, the second data storage device includes a memory configured for storing data and a control module coupled to the memory that controls the transmission of data from the memory to the second power integrated network and the storage of data received from the second power integrated network in the memory. The second data storage device further includes a power module coupled to the memory and the control module. The power module receives power from the second power integrated network to energize the second data storage device. The memory may be device including a SCSI interface.

In another embodiment, the second data storage device includes a network interface coupled to the control module where the network interface is configured to communicate with the second power integrated network. The control module may package data for transmission over the second power integrated network. In yet another embodiment, the power module includes a power converter for converting the power received from a first voltage level to a second voltage level. The second voltage level may be lower than the first voltage level.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing description of various embodiments of the invention should be appreciated more fully from the following further description thereof with reference to the accompanying drawings wherein:

Figure 1A is a schematic diagram of a prior art local area network ("LAN") including a storage device.

Figure 1B is a schematic diagram of a prior art storage area network ("SAN") coupled to a local area network.

Figure 2 is a schematic diagram of a data storage device in a power integrated network in accordance with an embodiment of the invention.

Figure 3 is a schematic diagram of a data storage device for use in a power integrated network in accordance with an embodiment of the invention.

Figure 4 is a schematic diagram of a redundant data storage system in accordance with an embodiment of the invention.

DESCRIPTION OF PREFERRED EMBODIMENTS

In an embodiment of the present invention, a data storage device is provided that may be connected to a power integrated network (*e.g.*, a local area network implementing Power Ethernet, discussed below). The data storage device includes a memory configured to store data and a control module for controlling the transfer of data between the memory and the power integrated network across the network interface. The data storage device also includes a power module for receiving and converting power from the power integrated network. The power received from the power integrated network is used to energize the components of the data storage device. Accordingly, the data storage device is able to receive both data and power from the power integrated network.

As used herein, the term "power integrated network" refers to a local network that transmits both power and data to member computer systems (e.g., network appliances, personal computers, or servers) in the network. Such power may or may not be used by the member computer systems. Among other ways, the power and data may be transmitted on a single cable via different wires, or via the same wire. Illustrative power integrated networks include Power Ethernet networks, which implement the proposed IEEE 802.3af standard. Currently, this standard is in draft form and is expected to be completed and adopted sometime in late 2001. Computer systems utilizing this standard are capable of receiving power (e.g., about fourteen watts) and data from a computer cable across an IEEE DTE (data terminal equipment) through a MDI (media dependent interface) compliant port. In addition, a Power Ethernet network may determine whether devices coupled to the network are capable of receiving power from the network.

It should be noted that the term "data" is used herein to broadly represent any type of information that may be electronically transmitted across a network. Such information may include, among other things, information commonly referred to as audio, video, signaling, control and data. In addition, instead of using the term "datum", the term "data" is used herein to represent both the singular and plural form of the term "datum."

Figure 2 is a schematic diagram of a power integrated network 206 that includes a number of network devices, *i.e.*, workstations 208, a server 204 and a data storage device 202. Server 204 may be any type of server used in a local area network, such as, for example, a stand-alone hardware server, or a server implemented as software on a router or a multiplexer. The power integrated network 206 is also coupled to a public network 210, *i.e.*, the Internet 210, in order to communicate with remote network devices in other networks. It should be noted that a different number of servers 204, workstations 208 and data

storage devices 202 than that shown in Figure 2 may be used in the power integrated network 206.

Data storage device 202 is advantageously configured so that it may be connected to the power integrated network 206 to receive both power and data rather than coupled to the power integrated network via another network device, such as server 204. Accordingly, data storage device 202 is configured to receive both data and power from the power integrated network 206. With this configuration, any-to-any connectivity may be provided between the data storage device 202 and the other network devices in the power integrated network 206.

Figure 3 is a schematic diagram of a data storage device in accordance with an embodiment of the invention. As discussed above, data storage device 302 is configured so that it may be connected to a power integrated network 306 to exchange data with the power integrated network 306 as well as receive power from the power integrated network 306.

Data storage device 302 includes, among other things, a network interface 308 to communicate data with the power integrated network 306. The network interface may be a physical interface such as, for example, a conventional plug for plugging into an Ethernet network or a computer card with the attendant functionality, and/or a software module that, when executing, performs the desired functions. When data is received by the data storage device, the network interface 308 forwards the data to a memory 312. Memory 312 is used for storing data. In embodiments of the invention, memory 312 may be any type of device used to store data, such as, for example, disk drives, tape drives, optical drives, RAMS's, other static memory devices and PCMCIA (Personal Computer Memory Card International Association) cards. In a further embodiment, memory 312 may be a device implementing the Small Computer System Interface (SCSI) standard.

A control module 310 is coupled to the memory 312 and the network interface 308. Control module 310 is used to control the transmission of data from the memory 312 to the power integrated network 306 and the storage of data received from the power integrated network 306 in the memory 312.

- 5 Control module 310 communicates with the power integrated network 306 via the network interface 308. Control module 310 may utilize protocols, known in the art, that permit communication between a SCSI device and a local area network, such as an Ethernet network. One such protocol is described in an IETF Internet draft document entitled "The SCSI Encapsulation Protocol (SEP)", which is referenced as draft-wilson-sep-00.txt (May 2000), and is herein incorporated by reference in its entirety. In illustrative embodiments, the control module is a software component executing on a microprocessor.

- 10 In one embodiment, control module 312 may also include registration logic (not shown) that is used to register with a central authority, such as a server, when the data storage device 302 is connected to the power integrated network 306. An exemplary registration process is disclosed in co-pending U.S. Patent Application Serial No. _____, attorney docket number 2204/A60, filed December 21, 2000, entitled "Apparatus and Method for Confirming Physical Connection of a Network Device," which is hereby incorporated by reference. The control module may also include other logic not shown in Figure 3. Control module 310 may include, for example, an IP stack for executing all network communication processes and multicast functionality.

- 20 The IP stack preferably includes the link layer, network layer and transport layer software and/or hardware that permits the data storage device to transmit to and receive data from the power integrated network. The multicast functionality may be implemented as a software device that implements one or more multicast protocols (e.g., Protocol Independent Multicast, also known as

"PIM"). The multicast functionality permits the data storage device to transmit data message to other network devices on the power integrated network.

As discussed above, data storage device 302 may be powered by the power integrated network 306. Accordingly, the data storage device 302 also includes a power module 304 that receives power from the power integrated network 306. The power received may be used to energize the components of the data storage device 302. Power module 304 may include a power converter (not shown) to convert the voltage level of the power received from the power integrated network 306 to a voltage level that may be used by the various components of the data storage device 302. In one embodiment, the power received is a constant DC power supply. Accordingly, the power converter may be a DC power circuit that adjusts the voltage received from the power integrated network 306 to an appropriate level for the data storage device 302 and its components. For example, the DC power circuit may be a down converter that converts an incoming DC voltage to a lower DC voltage level.

An application of the data storage device 302 is shown in Figure 4. Figure 4 is a schematic diagram of a redundant data storage system in accordance with an embodiment of the invention. Data storage devices 402A-D are each connected to a power integrated network 406 at a different physical location. Data storage devices 402A-D may be used in a RAID (Redundant Array of Independent Disks) configuration. In illustrative embodiments using a RAID 3 or RAID 5 configuration, for example, at least one of the data storage devices will include parity information. Each data storage device 402A-D may be connected to a power integrated network at a different physical location. For example, data storage device 402A may be connected to a network in New York; data storage device 402B may be connected to a network in Boston; data storage device 402C may be connected to a network in Los Angeles; and data storage device 402D may be connected to a network in Denver. A server 404 at one of the locations

may be used as a host system. The data storage devices 402 in the RAID storage system will appear as one single device to the server 404. The data storage devices 402A-D may communicate with server 404 via a public network, i.e., the Internet 410 that is coupled to each power integrated network 406 and hence, to each data storage device 402A-D. In the system as shown in Figure 4, if there is a physical failure at one location, the other data storage devices in the RAID array will not be affected by the failure. As mentioned above, each power integrated network 406 may be coupled to the Internet 410 to enable the data storage devices 402A-D and server 404 to communicate. The data storage devices 402A-D and server 404 may communicate over the Internet by means of a known protocol such as TCP/IP (Transmission Control Protocol with Internet Protocol).

Some aspects of the invention may be implemented at least in part in any conventional computer programming language. For example, some embodiments may be implemented in a procedural programming language (e.g., "c") or an object oriented programming language (e.g., "C++"). Alternative embodiments of the invention may be implemented as preprogrammed hardware elements (e.g., application specific integrated circuits, FPGAs, and digital signal processors), or other related components.

Thus, the present invention may be embodied as a data storage device for use in a power integrated network that includes a memory configured to store data and a control module coupled to the memory to control the transmission of data from the memory to the power integrated network and the storage of data received from the power integrated network in the memory. A power module coupled to the memory and the control module receives and converts power received from the power integrated network from a first voltage level to a second voltage level.

Embodiments of the present invention may be embodied as a method for providing data storage in a power integrated network that includes coupling a

data storage device to the power integrated network where the data storage device is configured to communicate with the power integrated network. The method also involve receiving data and power at the data storage device from the power integrated network and using the power received from the power
5 integrated network to energize the data storage device.

Embodiments of the present invention may also be embodied as a system for providing data storage in a communication network that includes a first power integrated network and a second power integrated network that provide data and power. The first power integrated network and the second power
10 integrated network are both coupled to a specified network so that data may be transmitted from and received by each power integrated network across the specified network. A first data storage device is coupled to the first power integrated network and is configured to communicate with the first power integrated network and to receive power from the second power integrated
15 network. A second data storage device is coupled to the second power integrated network and is configured to communicate with the second power integrated network and to receive power from the second power integrated network.

Although various exemplary embodiments of the invention have been
20 disclosed, it should be apparent to those skilled in the art that various changes and modifications can be made that will achieve some of the advantages of the invention without departing from the true scope of the invention. These and other obvious modifications are intended to be covered by the appended claims. It should be noted that although Power Ethernet and this IEEE standard are
25 discussed, various embodiments are not limited to such standard. Accordingly, various embodiments can be used on other types of networks that transmit both power and data.